arXiv:1702.04683v1 [cs.LG] 15 Feb 2017

Distributed deep learning on edge-devices: feasibility via adaptive compression

Corentin Hardy*†, Erwan Le Merrer* and Bruno Sericola†
*Technicolor †INRIA

***Abstract*—The state-of-the-art results provided by deep learning come at the price of an intensive use of computing resources. The leading frameworks (*e.g.*, TensorFlow) are executed on GPUs or on high-end servers in datacenters. On the other end, there is a proliferation of personal devices with possibly free CPU cycles. In this paper, we ask the following question: *Is distributed deep learning computation on WAN connected devices feasible, in spite of the traffic caused by learning tasks*? We show that such a setup rises some important challenges, most notably the ingress traffic that the servers hosting the up-to-date model have to sustain.**

In order to reduce this stress, we propose *AdaComp*, a new algorithm for compressing worker updates to the model on the server. Applicable to stochastic gradient descent based approaches, it combines efficient gradient selection and learning rate modulation. We then experiment and measure the impact of compression and device reliability on the accuracy of learned models. To do so, we leverage an emulator platform we developed, that embeds the TensorFlow code into Linux containers. We report a reduction of the total amount of data sent by workers to the server by two order of magnitude (*e.g.*, 191-fold reduction for a convolutional network on the MNIST dataset), when compared to the standard algorithm based on asynchronous stochastic gradient descent, while maintaining model accuracy.

I. Introduction

Machine learning methods, and in particular deep learning, are nowadays key components for building efficient applications and services. Deep learning recently permitted significant improvements over state-of-the-art techniques for building classification models for instance [12]. Its use spans over a large spectrum of applications, from face recognition in [21], to natural language processing with word2vec [19] and to video recommendation in YouTube [6].

Learning a model using a deep neural network (we denote DNN hereafter) requires a large amount of data, as the precision of that model directly depends on the quantity of examples it gets as input and the number of times it iterates over them. Typically, the last image recognition DNNs, such as [12] or [9], leverage very large datasets (like Imagenet [23]) during the learning phase; this leads to the processing of over 10TB of data. The direct consequence is the compute-intensive nature of running such approaches. The place of choice for running those methods is thus well provisioned datacenters, for the more intensive applications, or on dedicated and GPU-powered machines in other cases. In this context, recently introduced frameworks for learning models using DNNs are benchmarked in cloud environments (*e.g.*, TensorFlow [2] with GPU-enabled servers and 16Gbps network ports).

Meanwhile, the number of processing devices at the edge of the Internet keeps increasing in a steep manner. In this paper, we explore the possibility of leveraging edge-devices for those intensive tasks. This new paradigm rises significant feasibility questions. The dominant *parameter server* computing model, introduced by Google in 2012 [7], uses a set of *workers* for parallel processing, while a few central servers (denoted the *parameter server* (PS) hereafter for simplicity) are managing shared states modified by those workers. Workers frequently fetch the up-to-date model from the PS, make computation over the data they host, and then return gradient updates to the PS. Since DNN models are large (from thousands to billions parameters [24]), placing those worker tasks over edge-devices imply significant updates transfer over the Internet. The PS being in a central location (typically at a cloud provider), the question of inbound traffic is also crucial for pricing our proposal. Model learning is facing other problems such as device crashes and worker asynchrony.

To illustrate the feasibility question, we implement the largest distribution scenario considered in the TensorFlow paper [2], where 200 machines are collaborating to learn a model. We measure the aggregated traffic at the PS, and plot it on Figure 1.

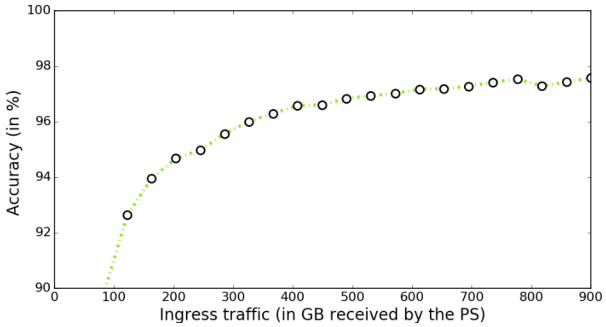

Fig. 1. Accuracy (in %) vs aggregated traffic (in GB) tradeoff: inbound aggregated ingress traffic (model updates) received at the Parameter Server (PS). For an Asynchronous Stochastic Gradient Descent (ASGD) with 200 workers, a convolutional neural network model training, on the MNIST dataset.

We observe a considerable amount of ingress traffic received by the PS, of the order a Terabyte for an accurate model. This amount of traffic is due to workers sending their updates to the PS, and is not even reported in research works, as well provisioned and dedicated data center networks are assumed.

Clearly, in a setup leveraging edge-devices, this amount of ingress traffic has to be drastically reduced to remove weight on both the Internet and on the PS. Meanwhile, the data to be processed at edge-devices themselves is not the limiting factor (3.9MB each in this experiment, as the dataset is split among workers). Our solution is to introduce a novel compression

technique for sending updates from workers to the PS, using gradient selection [25]. We thus study the model correctness with regards to update compression, as well as with regards to device reliability.

The main contributions of this paper are the following:

- Exposing the parameter server model implications (on ingress traffic at the PS and on model accuracy), in an edge-device setup.
- Introducing *AdaComp*, a novel compression technique for reducing the ingress traffic at the PS.
- Extensively experimenting *AdaComp* within TensorFlow, and comparing it to competitors with regards to model accuracy.

The remaining of this paper is organized as follows. Section II briefly presents the basics of deep learning, and of learning in datacenters. Section III introduces the considered execution setup on edge-devices, and recalls important performance metrics for deep learning. Section IV presents the *AdaComp* algorithm, for reducing the ingress traffic. We then evaluate *AdaComp* and its competitors in Section V, before presenting related works in Section VI and concluding the paper in Section VII.

II. Distributed Deep Learning and the Parameter Server Model

A. Basics on DNN training on a single core

Deep neural networks (or DNNs) are machine learning models used for supervised or unsupervised tasks. They are composed of a large succession of layers. Each layer l is the result of a non-linear transformation of the previous layer (or input data) given a weight matrix V_l and a bias vector β_l. In this paper, we focus on supervised learning, where a DNN is used to approximate a target function f (*e.g.*, a classification function for images).

For instance, given a training dataset $\mathcal{D}$ of images $x_h \in X$ and associated labels $y_h \in Y$, a DNN has to adapt its parametric function $\hat{f}_\Theta : X \to Y$, with Θ a vector containing the set of all parameters, *i.e.*, all entries of matrices V_l and vectors β_l. Training is performed by minimizing a loss function L which represents how good is the DNN at approximating f.

$$L(\Theta) = \sum_{(x_h, y_h) \in \mathcal{D}} e(y_h, \hat{f}_\Theta(x_h)),$$

where $e(y_h, \hat{f}_\Theta(x_h))$ is the error of the DNN output $\hat{f}_\Theta(x_h)$ for x_h with parameters Θ. The minimization of $L(\Theta)$ is performed by an iterative update process on the parameters called gradient descent (GD). At each step, the algorithm updates the parameters as follows (vector notations) :

$$\Theta(i+1) = \Theta(i) - \alpha \Delta\Theta(i),$$

with the learning rate $\alpha \in (0, 1]$ and $\Delta\Theta(i)$ approximating the gradient of the error function. $\Delta\Theta(i)$ is computed by a *back-propagation* step [14] on $\mathcal{D}$ at time i. In the following sections, we use a more efficient variant of GD, called Stochastic Gradient Descent (SGD), which processes only a mini-batch of the training dataset per iteration.

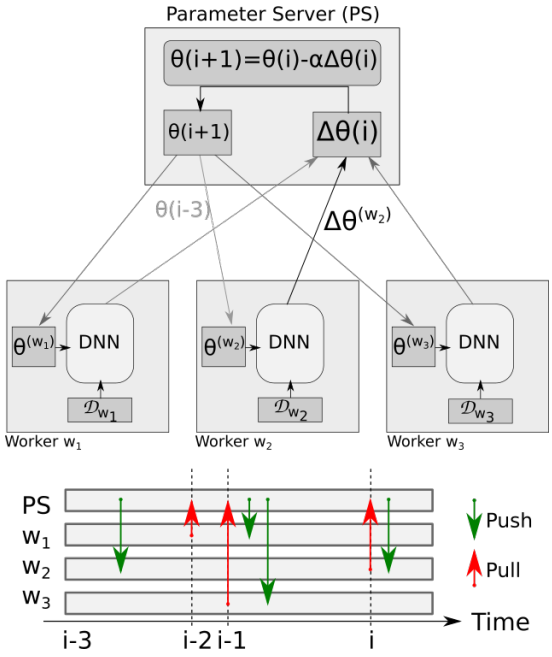

Fig. 2. The asynchronous *parameter server* computation model (with 3 workers) and an example of update chronology. Pull arrows represent vector refreshes from the PS to a worker, and Push arrows the updates sent by a worker to the PS.

B. Distributed Stochastic Gradient Descent in the Parameter Server Model

To speed-up DNN training, Dean et al. proposed in [7] the *parameter server* model, as a way to distribute computation onto up to hundreds of machines. The principal idea is to parallelize data computation: each compute node (or *worker*) processes a subset of the training data asynchronously. For convenience, we report the core notations in Table I.

The parameter server model is depicted on Figure 2. In a nutshell, the PS starts by initializing the vector of parameters to learn, Θ. Training data $\mathcal{D}$ is split across workers w_1, w_2, w_3. Each worker runs asynchronously, and whenever it is ready, takes the current version of the parameter vector from the PS, performs a SGD step and sends the vector update (that reflects the learning on its data) to the PS. In Figure 2, worker w_2 gets the state of Θ at iteration $i-3$, performs a SGD step and sends an update to the PS. Meanwhile, workers w_1 and w_3, which finish their step before w_2, send their updates. At each update reception, the PS updates Θ, and increments the timestamp of Θ (*e.g.*, $\Theta(i-3)$ becomes $\Theta(i-2)$ after the reception of w_1's update). An important parameter is the mini-batch size, we denote b, that corresponds to the subset size of the local data that each worker is computing upon, before sending one update to the PS (*e.g.*, if $b = 10$, an update is sent by w_1 after its DNN has processed 10 images of $\mathcal{D}_{w_1}$).

Note that in the PS model, the fact that n workers are training their local vector in parallel and then send the updates introduces concurrency, also known as *staleness* [28]. The staleness of the local vector $\Theta^{(w)}$ for a worker w is the number of times Θ has evolved in between the fetch by w, and the time where it itself pushes an update $\Delta\Theta^{(w)}$ (*e.g.*, on Figure 2, W_2's update staleness is equal to 2).

Notation	Meaning
PS	The parameter server, hosting the up-to-date parameter vector Θ
n	The number of workers (excluding the PS)
$w_1, \ldots, w_n$	The identifiers of the n workers
b	Size of mini-batches: the number of items processed by a worker before sending an update
α	The model learning rate
c	The compression ratio (% of selected parameters [per matrix V_l or vector β_l])
$\mathcal{D}_w$	Training data local to a worker w. Full dataset is thus $\mathcal{D} = \cup_{i=1}^{n} D_{w_i}$
I	The number of performed iterations (updates received and applied by the PS)
$\Theta(i)$	Vector containing all parameters at iteration i, *i.e.*, all entries of matrices V_l and vectors β_l at iteration i
$\Theta_k(i)$	The k-th entry of $\Theta(i)$
$\Delta\Theta(i)$	Vector containing all updates to Θ
$\sigma(i)$	The global staleness of update i
$\sigma_k(i)$	Local staleness of an update associated to the parameter Θ_k
$\alpha_k(i)$	Local learning rate associated to the parameter Θ_k given $\sigma_k(i)$
$\Theta^{(w)}$, $\Delta\Theta^{(w)}$ and $\Delta\tilde{\Theta}^{(w)}$	The local parameter vector state at worker w, the update computed from $\Theta^{(w)}$ by worker w and the resulting compressed update

TABLE I. LIST OF NOTATIONS, FOR A DISTRIBUTED DEEP LEARNING SETUP.

III. DISTRIBUTED DEEP LEARNING ON EDGE-DEVICES

A. Execution setup

The execution setup we consider replaces the datacenter worker nodes by edge-device nodes (*e.g.*, personal computers of volunteers with SETI@home, or home gateways [27]) and the datacenter network by the Internet. The PS remains at a central location. In this context, we assume that the training data reside with the workers; this serves for instance as a basis for privacy-preserving scenarios [25], where users have their photos at home, and want to contribute to the computing of a global photo classification model, but without sending their personal data to a cloud service.

The formal execution model remains identical to the PS model: workers have enough memory to host a copy of Θ, data is shuffled on the n workers, and finally they agree to collaborate to the machine learning task (tolerance to malicious behaviours is out of the scope of this paper). The stringent aspect of our setup is the lower connectivity capacity of workers, and their expected smaller reliability [27]. Regarding connectivity, we have in mind a standard ASDL/cable setup where the bottleneck at the device is the uplink (with *e.g.*, respectively 100Mb/10Mb for down/up bandwidth); we thus optimize on the worker upload capacity by compressing the updates it has to send *i.e.*, the push operation, rather than the opposite *i.e.*, the pull operation.

B. Staleness mitigation

Asynchronous vector fetches and updates by workers involve perturbations on the SGD due to the staleness of local $\Theta^{(w)}$. The staleness, proportional to the number of workers (please refer to [28], [16] or [8] for in depth phenomenon explanation), is an important issue in asynchronous SGD in general. A high staleness is to be avoided, as workers that are relatively slow will contribute, through their update, to a Θ that has evolved (possibly significantly) since they last fetched their copy of Θ. This factor is known for slowing down the computation convergence [28]. In order to cope with it, works [28] and [20] propose to adapt the learning rate α as a function of the current staleness to reduce the impact of stale updates, which will also reduce the number of updates needed to train Θ.

C. Reducing ingress traffic at the PS

In the edge-device setting, where devices collaborate to the computation of a global Θ, we argue that the critical metric is the bandwidth required by the PS to cope with incoming worker updates. This allows for frontend server/network dimensioning, and is also a crucial metric, as cloud providers often bill on upload capacities to cloud servers (see *e.g.*, Amazon Kinesis, which charges depending on the number of 1MB/s message queues to align as frontends). Since our setup implies best effort computation at devices, not to saturate uplink, the workers are sending updates to the PS as background tasks. We thus measure the total ingress traffic at the PS collection point, in order to have an aggregated view of the uploaded traffic (thus from workers) incurred by the deep learning tasks.

a) Mini-batch size: In a datacenter, the aim of distributing a deep learning task is to get processing speed-up, as compared to computation on a single machine. Gupta *et al.* in [8] show that a good precision is achievable with a given number of *epochs*[1] by keeping $n \times b$ constant (with n the number of workers and b the mini-batch size). In other terms, this means that if we increase n, we should reduce b to maximize the speed-up. But reducing b increases the number of updates needed to train Θ, and thus increases the total traffic on the network. Our proposal does not preserve the $n \times b$ ratio, thus introducing more computation for workers but in turn reducing updates.

b) Compression: Shokri *et al.* [25] proposed a compression mechanism for reducing the size of the updates sent from each worker to the PS, named *Selective Stochastic Gradient Descent*. At the end of an iteration over its local data, a worker sends only a subset of computed gradients, rather than all of them. The selection is made either randomly, or by keeping only the largest gradients. The compression ratio is represented by fixed value $c \in (0, 1]$. They experimentally show that model accuracy is not impacted much by this compression method for a SGD in a single core (with *e.g.*, a MLP model accuracy

[1] One *epoch* corresponds to the whole dataset having being processed once.

decreasing from 98,10% to 97,07% accuracy, on the MNIST dataset with a selection of only 1% of parameters per update).

In the light of this Section, the total amount of ingress traffic received at the PS is in the order of $I \times M \times c$, with M the size of Θ (*e.g.*, in MB). As the crucial focus of machine learning is the reached model accuracy, we saw (Figure 1) that unfortunately it is not linear with the amount of data received at the PS. That is why we measure the accuracy/ingress traffic tradeoff experimentally, in the remaining of this paper. As we shall see in the evaluation Section, the use of approach [25] manages to reduce the size of updates, but at the cost of accuracy. This is a clear impediment, because deep learning is leveraged for its state-of-art accuracy results. In order to cope with both ingress traffic and accuracy maintenance, we now introduce the *AdaComp* algorithm.

IV. Compressed updates with *AdaComp*

AdaComp is a solution to conveniently combine the two concepts of compression and staleness mitigation, for further compressing worker results. This permits drastic ingress traffic reduction, for a comparable accuracy with best related work approaches.

a) Rationale: To do so, we propose the following approach. We first observe that the content of updates pushed by workers are often *sparse* (*i.e.*, most of the parameters have not changed, $\Delta\Theta_k \approx 0$ for most of Θ_k): a selection method based on largest values of update is a good solution to compress it with little loss of information. Second, we observe that staleness mitigation is handled solely at the granularity of a whole update, in related approaches. From those remarks, we propose in *AdaComp* to compute staleness not on an update, but *per parameter*. Our intuition is that sparsity and staleness mitigation on individual parameters will allow for increased efficiency in asynchronous worker operation, by removing a significant part of update conflicts. An independent staleness score computed for each parameters is a reasonable assumption: F. Niu *et al.* [22] consider parameters as independent during the SGD process.

b) Selection method: We propose a novel method for selecting gradient at each worker: only a fraction c of the largest gradient per matrix V_l and vector β_l are kept. This selection permits to balance learning across DNN layers and better reflects computed gradients as compared to a random selection (or solely the largest ones across the whole model) as in [25].

c) Algorithm details: AdaComp operates in the following way, also described with pseudo-code in Algorithm 1 and 2. The PS keeps a trace of all received updates at given timestamps (as in [8], [28]). When a worker pulls Θ, it receives the associated timestamp j. It computes a SGD step and uses the selection method to compress the update. Meanwhile, intermediate updates (pushed by other workers) increase the timestamp of the PS to $i \geq j$. Instead of computing the same staleness for each $\Delta\Theta_k(i)$ of the update $\Delta\Theta(i)$, we define an adaptive staleness for each parameter Θ_k as follows :

$$\sigma_k(i) = \sum_{u=j}^{i-1} \mathbb{1}_{\{\Delta\Theta_k(u)\neq 0\}}, \tag{1}$$

where $\mathbb{1}_A$ is the indicator function of condition A equal to 1 if condition A is true and 0 otherwise. The staleness σ_k is then computed **individually** for each parameter by counting the number of updates applied on it since the last pull by the worker. We use the update equation inspired by [28]:

$$\Theta_k(i+1) = \Theta_k(i) - \alpha_k(i)\Delta\Theta_k(i), \tag{2}$$

where

$$\alpha_k(i) = \begin{cases} \alpha/\sigma_k(i) & \text{if } \sigma_k(i) \neq 0 \\ \alpha & \text{otherwise.} \end{cases} \tag{3}$$

The α_k is the learning rate computed for the k-th parameter of Θ given staleness σ_k. The parameter Θ_k that was updated since the worker pull, will see α_k reduced as a function of the staleness, as shown in equation (3). This method has the effect of reducing the concurrency between the possibly numerous asynchronous updates taking place along the learning task, while tacking advantage of gradient selection to reduce update size[2].

Algorithm 1 *AdaComp* at workers

procedure WORKER($PS, \mathcal{D}_w, I, b, c$)
  $i \leftarrow 0$
  while $i < I$ **do**
    $\Theta^{(w)}, i \leftarrow$ Pull_$\Theta(PS)$ ▷ get current $\Theta(i)$ from PS
    $\Delta\Theta^{(w)} \leftarrow$ SGD_STEP(D_w, b)
    $\Delta\tilde{\Theta}^{(w)} \leftarrow$ SELECT_GRAD($\Delta\Theta^{(w)}, c$)
    Push($\Delta\tilde{\Theta}^{(w)}, i$) ▷ update push to PS, with fetched i
  end while
end procedure
procedure SGD_STEP($\mathcal{D}_w, b$)
  Select a mini-batch B by sampling b examples from $\mathcal{D}_w$ and compute $\Delta\Theta^{(w)}$ with back-propagation method.
  return $\Delta\Theta^{(w)}$
end procedure
procedure SELECT_GRAD($\Delta\Theta^{(w)}, c$)
  Select a subset $\Delta\tilde{\Theta}^{(w)}$ of $\Delta\Theta^{(w)}$ by keeping the $(100\times c)\%$ largest parameters, in absolute value, of each matrix V_l and vector β_l.
  return $\Delta\tilde{\Theta}^{(w)}$
end procedure

V. Experimental Evaluation

In this section, we present the results obtained by *AdaComp* on two different DNNs, and compare them to those of two competitors (described in section V-B).

[2]Please note that the best low level data representation for update encoding is out of the scope of this paper: it will causes at best a reduction of a small factor w.r.t. python serialization (*i.e.*, the TensorFlow language), while we target in this paper two orders of magnitude decrease on the network footprint by targeting the algorithmic scope of distributed deep learning.

Algorithm 2 *AdaComp* at the parameter server

procedure PS(α, I)
  Initialize $\Theta(0)$ with random values.
  for $i \leftarrow 0, I$ **do**
    Get($\Delta\tilde{\Theta}^{(w)}, j$) ▷ wait a push
    $\Delta\Theta(i) \leftarrow \Delta\tilde{\Theta}^{(w)}$
    for all $\Delta\Theta_k(i) \in \Delta\Theta(i)$ **do**
      $\sigma_k \leftarrow \sum_{u=j}^{i-1} \mathbb{1}_{\{\Delta\Theta_k(u)\neq 0\}}$
      if $\sigma_k = 0$ **then** $\alpha_k \leftarrow \alpha$ **else** $\alpha_k \leftarrow \alpha/\sigma_k$
      $\Theta_k(i+1) \leftarrow \Theta_k(i) - \alpha_k \Delta\Theta_k(i)$
    end for
  end for
end procedure

A. Experimental platform

For assessing the value of *AdaComp* in a controlled and monitored environment, we choose to emulate the overall system on a single powerful server. The server allows us to remove the hardware and configuration constraints. We choose to represent edge-devices through Linux containers (LXC) on a Debian high end server (Intel(R) Xeon(R) CPU E5-2667 v3 @ 3.20GHz CPUs, a total of 32 cores, and 1/2TB of RAM). Each of them runs a TensorFlow session to train locally the DNN[3]. The traffic between LXCs can then be managed by the host machine with *virtual Ethernet* (or veth) connections.

An experiment on this platform is as follows. A set of n LXC containers is deployed, to represent the workers. Each worker in a container has access to a proportion $1/n$ of the training dataset. One LXC container is deployed to run the PS code. All workers are connected to PS by a veth virtual network. Finally, one last LXC container is deployed to evaluate the accuracy evolution of Θ with the test dataset.

During the execution, the platform adds a random waiting time, for each worker, between the effective computation time and the push step to the PS. This time ensures that the order of worker updates is random, mimicking the possible variety in hardware or resource available to the workers. This random draw also allows us to ensure that the PS container is not overloaded by worker updates, during the emulation on our server. We report a runtime of about $1,5$ day for each run (for reaching $I = 250,000$ iterations).

In such a setup, the interesting metric to observe is the accuracy reached according the number of iterations performed, which is the same, whether computation takes place in an emulated setup or in a real deployment. The benefit of our platform is a tight monitoring of the resulting TCP traffic, during the learning task.

B. Experimental setup, metrics and competitors

We experiment our setup with the MNIST dataset [15], also used by our competitor [25]. The goal is to build an image classifier able to recognize handwritten digits (*i.e.*, 10 classes). The dataset is composed of a training dataset with $60,000$ images, *i.e.*, $6,000$ per class, and a test dataset with $10,000$ images. Each image is represented by by 28x28 pixel with a 8bit grey-level matrix.

The experiments are launching $n = 200$ workers, and one PS, corresponding to the scale of operation reported by the TensorFlow paper [2].

1) Learning with two models: We test two different DNN models to build the classifier. The first one is a Multi-Layers Perceptron (MLP) containing three fully connected layers ($109,386$ parameters), and is taken from the Shokri et al. paper [25]. The second is a Convolutional Neural Network (CNN), consisting of two convolutional layers and two full-connected layers ($211,690$ parameters), and is taken from the Keras library [5] for the MNIST dataset.

2) Competitors: We compare the performance of *AdaComp* to the ones of 1) the basic Async-SGD method (which we denote ASGD) [16] as a baseline. This algorithm maintains accuracy based on a fix $n \times b$ ratio as discussed in Section III-C. 2) the *Comp-ASGD* algorithm, which is similar to ASGD, but implements a gradient selection as described in [25]. For both of competitor, the learning rate is divided by staleness is and global to each update.

3) Evaluation metrics: We measure the performance in terms of 1) accuracy, according the number of iterations i, using the MNIST test dataset (expressed as a percentage of correct classifications using the learned $\Theta(i)$), and 2) the caused ingress traffic at the PS. We then record the communication traffic between the workers and the PS using standard Unix tools, such as *ifconfig*. For *AdaComp*, we finally highlight the evolution of gradients during the learning phase.

C. Accuracy results

All experiments are executed until the PS has received a total of $I = 250,000$ updates from the workers; this corresponds to an upper bound for convergence of the best performing approaches we present in this Section. Most of the experiments have been run 3 times; plots show average, minimal and maximal values of the moving average accuracies for each run. Final test accuracy for an experiment is the mean of the maximal accuracy reached by each run.

1) MLP: The first experiment consists in training the MLP model to classify the MNIST digits.

Figure 3 (left) plots the accuracy of the model depending of the number of iteration. Figure 3 (right) reports the accuracy as function of the ingress traffic, measured at the PS. ASGD, *Comp-ASGD* and *AdaComp* are experimented within the same setup. For ASGD, we use two different batch sizes : $b = 1$ (representing the baseline), and $b = 10$ (the basic value we adopt for all other competitors). For both *Comp-ASGD* and *AdaComp*, we test the compression ratios $c = 0.1$ and $c = 0.01$, respectively representing 10% and 1% of the local parameters sent as an update by a worker to the PS.

The ASGD algorithm ($b = 1$) reaches a 95.27% accuracy at the end of the experiment; an important part of the final error is due to asynchronous updates (that are frequent in this setup with $n = 200$). ASGD with $b = 10$ has a slightly higher accuracy of 95.43%. Notably, the compression method used in *Comp-ASGD* causes accuracy results (94.80% and 92.10%

[3]Note that TensorFlow is also running in containers, while executed in a datacenter environment [1].

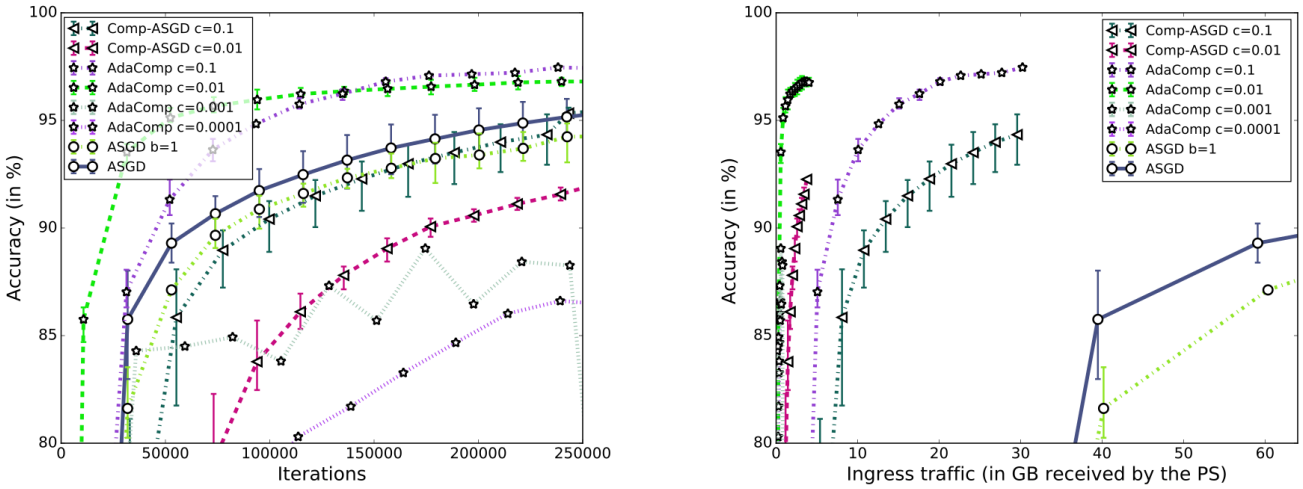

Fig. 3. MLP model learning on the edge-device setup: accuracy results (Left), and resulting ingress traffic (Right). Plotted results for *AdaComp* and its two competitors.

respectively for $c = 0.1$ and 0.01)) to be lower than the ASGD baseline. *AdaComp* is able to beat both ASGD and *Comp-ASGD* by a large margin (with 97.61% and 97.44%). To better understand the influence of the compression ratio on the accuracy of *AdaComp*, we decrease c to 0.001 and 0.0001. The resulting accuracy is significantly lower (with 91.48% and 87.62% respectively). We conclude that $c = 0.1$ is the sweet spot for *AdaComp* in this setup.

Regarding the resulting ingress traffic by considered algorithms, presented on Figure 3 (right), we observe striking differences. As expected, not relying on compression causes ASGD to produce a large amount of traffic in destination to the PS (up to 600GB after $250,000$ iterations), while the accuracy is still not at its possible best. The effect of compression for *Comp-ASGD* and *AdaComp*, with low values of c, is clearly noticeable. This allows to drastically reduce the amounts of traffic the PS has to deal with, and then in turn reduces the operational cost of the deep learning task.

Experiment conclusion To better compare the three competitors, we fix an accuracy level of 94%, and report the resulting aggregated ingress traffic for each algorithm. *AdaComp* generates 0.38GB of ingress traffic (with $c = 0.001$), representing 251 times less traffic than the best performing ASGD experiment ($b = 10$). Meanwhile, *Comp-ASGD* manages to reach this accuracy level for only 2 runs, with $c = 0.1$; it generates 19.85GB of ingress traffic.

2) CNN: We perform the next experiment using a CNN model, that is known to result in a better accuracy on image classification [13], [12], [26].

Figure 4 (left) first shows that the accuracy trends of the three experimented algorithms remain the same than for the MLP model. Compression also affects *Comp-ASGD* results, while it does not prevent *AdaComp* to perform better than both ASGD and *Comp-ASGD*. Final accuracy results are for ASGD 97.19% and 97.85% (with $b = 1$ and $b = 10$ respectively), for *Comp-ASGD* 96.07% and 95.11% ($c = 0.1$ and 0.01), and finally for *AdaComp* 98.7% and 98.59%.

Figure 4 (right) also shows a strictly identical trend for the PS ingress traffic produced by all three algorithms.

Experiment conclusion For this CNN model, and for an accuracy of 97%, *AdaComp* generates 1.33GB of ingress traffic, representing 191 times less traffic than ASGD ($b = 10$). Consistently with the MLP experiments, *Comp-ASGD* does not manage to reach this accuracy level.

3) Conclusion on the accuracy/ingress traffic tradeoff: *AdaComp* outperforms ASGD and *Comp-ASGD* on both accuracy and ingress traffic. While compressing the size of updates naturally leads to a reduced traffic flow to the PS, the fact that *AdaComp* beats ASGD is less intuitive. *AdaComp* counters the effects of staleness introduced by the asynchronous updates, by using fine grained updates. Indeed the probability of update conflicts of a parameter Θ_k is less important obviously than at the level of the whole set of parameters Θ_k in the update (as performed in *Comp-ASGD*). More precisely, *AdaComp* leverages a sparse selection of gradients, to allow a higher learning rate α_k for parameters not impacted by staleness (as described in Section IV).

This directly results in the ability for *AdaComp* to propose a better accuracy/ingress traffic tradeoff. The traffic needed when using *AdaComp*, in the order of GBs, is for instance compatible with edge-devices and background execution (*e.g.*, for a home gateway connected 24/7, this corresponds to an average of less than 5KB/s of data upload over a week). This makes it possible to consider the exploitation of edge-devices for learning models, while dimensioning the PS infrastructure as small as possible.

D. AdaComp *accuracy facing worker failures*

A key element for the adaption of distributed algorithms over unreliable resources is the capacity of those algorithms to provide accurate results despite failures. This experiment is thus targeted at the more unreliable nature of edge-devices that are used as workers in our setup. In order to accommodate other applications that are concurrently running in the edge-device the deep learning tasks may occur in background

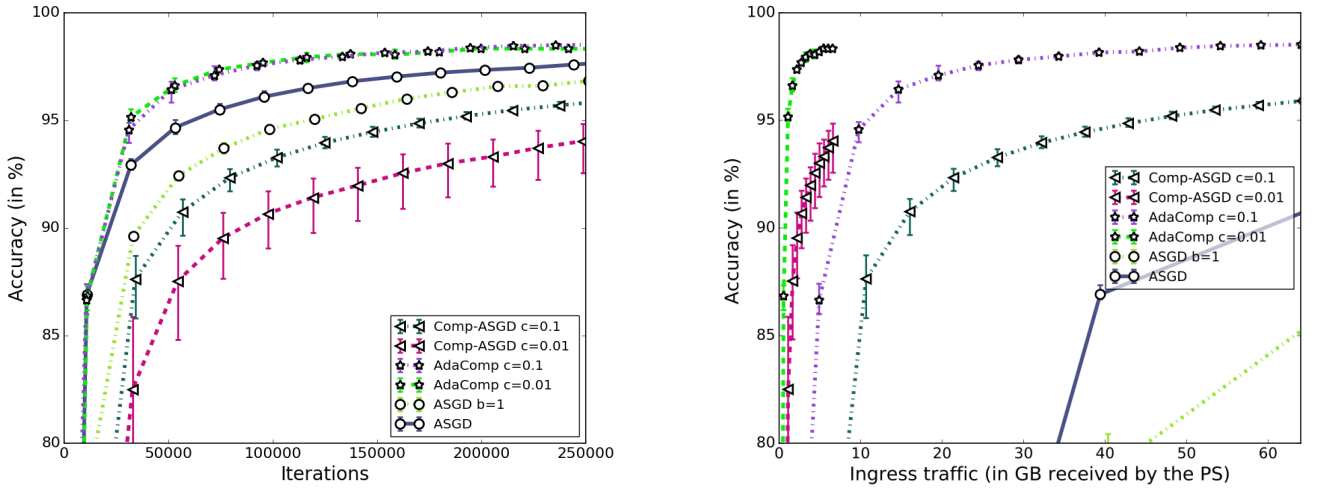

Fig. 4. CNN model learning in the edge-device setup: accuracy results (Left), and resulting ingress traffic (Right). Plotted results for *AdaComp* and its two competitors.

priority, and therefore need more time to complete. This is why we now consider fail-stop type of failures (*e.g.*, without warning messages or partial updates to the PS), where workers present at the beginning of the learning task may crash (with no reappearance). In addition, we consider that the local data of a crashed node are lost for the system.

We tune the probability p for each worker to crash after each push operation to $p = 0.0004$. On our platform, we simply freeze the container that has been randomly selected to act as crashed. In expectation, half of the initial population of 200 workers will have crash by the end of the experiment (*i.e.*, 200 workers are present at $I = 0$, and only 100 survived at $I = 250,000$). Results for *AdaComp* are presented on Figure 5. We operate on the MLP model.

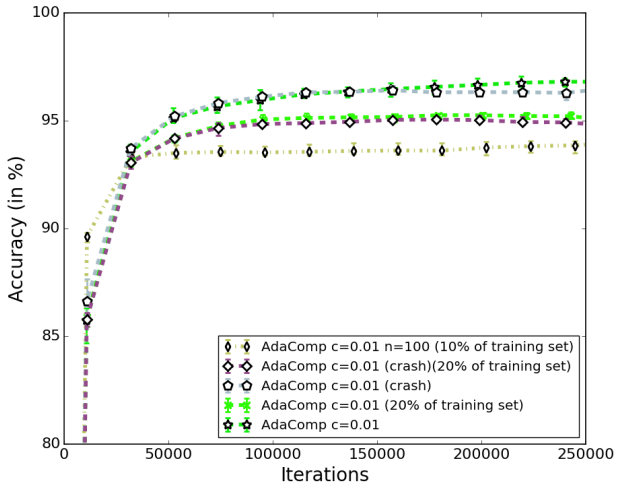
Accuracy (in %)
Iterations
AdaComp c=0.01 n=100 (10% of training set)
AdaComp c=0.01 (crash)(20% of training set)
AdaComp c=0.01 (crash)
AdaComp c=0.01 (20% of training set)
AdaComp c=0.01

Fig. 5. Accuracy of computed Θ, with half the workers crashing during the experiment (under a fail-stop crash model).

The first observation, with *AdaComp* $c = 0.01$ (crash) curve, is that crashes have very little effect on the final accuracy (to be compared to the *AdaComp* $c = 0.01$, where no crash occurs), with a final accuracy of 97.17% (*i.e.*, 0.27% less). This surprising result comes from the fact that the MNIST dataset is "too rich" for the learning task; in other words, we can reach a similar level of accuracy without learning on the whole dataset. This translates by nodes crashing with data that are not mandatory for the correctness of Θ in the end.

In order to observe the impact of failures on a "poorer" dataset, we run the same experiment with only 20% of the original MNIST dataset as training set (*i.e.*, 12000 images). We report a loss of accuracy of 0.23% between *AdaComp* with crashes and *AdaComp* with no crashes (with $c = 0.01$ for both), close to the previous loss of accuracy with the whole training dataset. To further explain this phenomenon, we run an experiment where only 100 workers participate to the learning task. We plot the *AdaComp* curve for $c = 0.01$, $n = 100$, 10% of the training set and no crash (*i.e.*, equivalent to a scenario where half of the 200 nodes would have disappeared with thus 10% of data prior to the start of the learning task). Final accuracy result is 94.35%, versus 95.24% for the previous experiment with crashes. This underlines that the 100 crashed nodes have in fact participated to the model learned, despite the fact that they have not computed over their whole local dataset (otherwise curve *AdaComp* with $c = 0.01(crash)(20\%$ of training set) would also have terminated around 94.35% accuracy as well).

Experiment conclusion This experiment underlines that crash failures of edge-devices will not affect the accuracy of the model if the dataset over which they learn is rich enough (thus introduces a form of data-redundancy). In the case where data is scarce, we note that the impact of failures remains very limited, provided that the workers have contributed at least a little before failing. Those two remarks are important for the adoption of distributed deep learning tasks, for they allow to consider operation on top of low-end and crash prone devices.

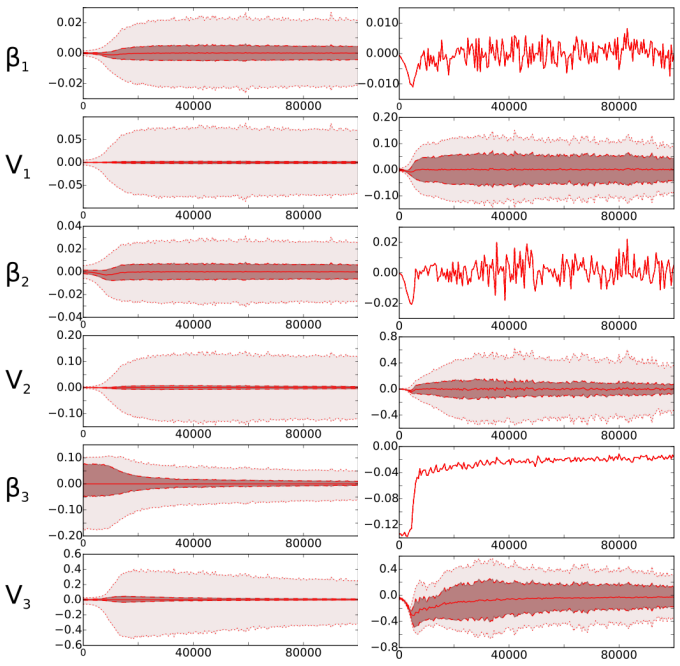

Fig. 6. The evolution of gradients during the first $100,000$ updates to Θ, for both ASGD with $b = 10$ (left) and for *AdaComp* with $c = 0.01$ and $b = 10$ (right). The plain line represents the average (or the sole value selected in the case of *AdaComp*). The outer lines represent $\min$ and $\max$ values, while the inner lines represent the positive and negative quartiles of those values (then of 50% of values in total).

E. *Visualization of the* AdaComp *learning phase*

To better understand the impact of the *AdaComp* algorithm on the model, we plot the distribution of gradients sent by workers during the learning phase (iterations) on Figure 6 (left), for the MLP model. We also plot the distribution for ASGD (rightmost figure), for comparison purposes. *AdaComp* selects the largest values for each matrix and vector, so its gradient distribution is more spread than that of the ASGD method for weight matrix V_i. For bias vector β_i, the compression is such that only one parameter is selected per update. A second interesting observation is that negative gradients are larger than positive one at the last layer; this leads *AdaComp* to mostly select them (and more specifically for bias).

VI. Related Work

The parameter server model is popular for distributing SGD with *e.g.*, DistBelief [7], Adam [4], and TensorFlow [2] as the principal ones. Increasing the number of workers in such an asynchronous environment causes staleness problems, that has been addressed either *1)* algorithmically, or by means of *2)* synchronization.

1) In [28], W. Zhang et *al.* propose to adapt the learning rate according to the current staleness, for each new update. They divide the learning rate by the staleness to limit the impact of highly stale updates. A recent paper of Odena [20] proposes to maintain the averaged variance of each parameters to weight new updates more precisely.

2) A simple solution to avoid staleness problems is to synchronize worker updates. However, waiting for all workers at each iteration is time consuming. W. Zhang et al. [28] propose the n-softSync method where the PS waits a fraction $1/n$ of all workers at each iteration. A more accurate update is computed by averaging their updates before applying them to the model. Another recent work of Chen *et al.* [3] shows that a synchronous SGD could be more efficient if the PS does not wait the last k workers at each iteration. In our setup, workers are user-devices without any guarantee on the upper bound of the response time. This calls for efficient asynchronous methods like *AdaComp*.

At the other extreme of the parallelism versus communication tradeoff, so-called *federated optimization* has been introduced [11], [10], [18]. A model is learned given a very large number of edge-devices each only processing over a few data, and each being equipped with a poor connection. A subset of active workers (which change at each iteration) run many iterations over their data before sharing their model. Global updates are then performed synchronously. Federated learning is thus communication-efficient, but does not take advantage of data parallelism for speeding-up computation, which is the goal of the parameter server model and of our *AdaComp* algorithm.

Researchers have also been interested in the distribution of other applications over edge-devices. This is for instance the case with so called *nano data centers* [27], where home gateways are the devices chosen for best effort general computation tasks. The performance focus of this work is on energy savings, as compared to a datacenter setup. Another example is the Wuala data storage service, that stores part of user data in a peer-to-peer fashion on edge-devices, to reduce the storage operational costs of the service provider [17]. Our distributed setup can also reduce operational cost if the users are willing to contribute to the model training, for instance in exchange of enhanced privacy for their content (their data is not uploaded to the cloud for processing), and may also reduce energy needs when running on devices such as home gateways, that are already up 24/7.

VII. Conclusion

We discussed in this paper the most important implications of running distributed deep learning training tasks on edge-devices. Because of upload constraints of devices and of their lower reliability in such a setup, asynchronous SGD is a natural solution to perform learning tasks; yet we highlighted that the amounts of traffic that have to transit over the Internet are considerable, in particular the PS ingress traffic, if no action is taken to adapt algorithms.

We proposed *AdaComp*, a new algorithm to compress updates by adapting to their individual parameter staleness. We show that this translates into a 251 or 191-fold reduction of the ingress traffic at the parameter server, as compared to the asynchronous SGD algorithm (for respectively the MLP and CNN models on the MNIST dataset), and for an as well better accuracy. We believe that this large reduction of the ingress traffic makes it possible to consider the actual deployment of learning tasks on edge-devices, and will particularly benefits the dimensioning of the infrastructure for receiving device updates, at the parameter server.

Future works include the removal of the PS, for further distribution of the DNN learning task. The PS has by definition a central position. Having a fully distributed learning system would remove the need for several high-end servers in the cloud. Yet some crucial questions about model accuracy in such a setup have to be explored.

REFERENCES

[1] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. J. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Józefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. G. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. A. Tucker, V. Vanhoucke, V. Vasudevan, F. B. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng. Tensorflow: Large-scale machine learning on heterogeneous distributed systems. *CoRR*, abs/1603.04467, 2016.

[2] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard, M. Kudlur, J. Levenberg, R. Monga, S. Moore, D. G. Murray, B. Steiner, P. Tucker, V. Vasudevan, P. Warden, M. Wicke, Y. Yu, and X. Zheng. Tensorflow: A system for large-scale machine learning. In *OSDI*, 2016.

[3] J. Chen, R. Monga, S. Bengio, and R. Jozefowicz. Revisiting distributed synchronous SGD. In *International Conference on Learning Representations Workshop Track*, 2016.

[4] T. Chilimbi, Y. Suzue, J. Apacible, and K. Kalyanaraman. Project adam: Building an efficient and scalable deep learning training system. In *OSDI*, 2014.

[5] F. Chollet. Keras. https://github.com/fchollet/keras, 2015.

[6] P. Covington, J. Adams, and E. Sargin. Deep neural networks for youtube recommendations. In *Recys*, 2016.

[7] J. Dean, G. Corrado, R. Monga, K. Chen, M. Devin, M. Mao, M. aurelio Ranzato, A. Senior, P. Tucker, K. Yang, Q. V. Le, and A. Y. Ng. Large scale distributed deep networks. In *NIPS*, 2012.

[8] S. Gupta, W. Zhang, and F. Wang. Model accuracy and runtime tradeoff in distributed deep learning: A systematic study. In *ICDM*, Dec 2016.

[9] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, 2016.

[10] J. Konecný, B. McMahan, and D. Ramage. Federated optimization: Distributed optimization beyond the datacenter. *CoRR*, abs/1511.03575, 2015.

[11] J. Konečný, H. Brendan McMahan, F. X. Yu, P. Richtárik, A. Theertha Suresh, and D. Bacon. Federated Learning: Strategies for Improving Communication Efficiency. *CoRR*, abs/1610.05492, Oct. 2016.

[12] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*. 2012.

[13] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, November 1998.

[14] Y. LeCun, L. Bottou, G. Orr, and K. Muller. Efficient backprop. In G. Orr and M. K., editors, *Neural Networks: Tricks of the trade*. Springer, 1998.

[15] Y. LeCun, C. Cortes, and C. J. Burges. The mnist database of handwritten digits, 1998.

[16] X. Lian, Y. Huang, Y. Li, and J. Liu. Asynchronous parallel stochastic gradient for nonconvex optimization. In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *NIPS*. Curran Associates, Inc., 2015.

[17] T. Mager, E. Biersack, and P. Michiardi. A measurement study of the wuala on-line storage service. In *P2P*, 2012.

[18] H. B. McMahan, E. Moore, D. Ramage, and B. A. y Arcas. Federated learning of deep networks using model averaging. *CoRR*, abs/1602.05629, 2016.

[19] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean. Distributed representations of words and phrases and their compositionality. In *NIPS*. 2013.

[20] A. Odena. Faster Asynchronous SGD. *ArXiv e-prints*, Jan. 2016.

[21] O. M. Parkhi, A. Vedaldi, and A. Zisserman. Deep face recognition. In *BMVC*, 2015.

[22] B. Recht, C. Re, S. Wright, and F. Niu. Hogwild: A lock-free approach to parallelizing stochastic gradient descent. In *NIPS*, 2011.

[23] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *IJCV*, 115(3):211–252, 2015.

[24] S. Shi, Q. Wang, P. Xu, and X. Chu. Benchmarking state-of-the-art deep learning software tools. *CoRR*, abs/1608.07249, 2016.

[25] R. Shokri and V. Shmatikov. Privacy-preserving deep learning. In *CCS*, 2015.

[26] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *CVPR*, 2015.

[27] V. Valancius, N. Laoutaris, L. Massoulié, C. Diot, and P. Rodriguez. Greening the Internet with Nano Data Centers. In *CoNext*, 2009.

[28] W. Zhang, S. Gupta, X. Lian, and J. Liu. Staleness-aware async-sgd for distributed deep learning. *CoRR*, abs/1511.05950, 2015.